arXiv:1702.02107v1 [cs.IR] 18 Jan 2017

First Study on Data Readiness Level

Hui Guan* Thanos Gentimis† Hamid Krim* James Keiser‡

Abstract

We introduce the idea of Data Readiness Level (DRL) to measure the relative richness of data to answer specific questions often encountered by data scientists. We first approach the problem in its full generality explaining its desired mathematical properties and applications and then we propose and study two DRL metrics. Specifically, we define DRL as a function of at least four properties of data: Noisiness, Believability, Relevance, and Coherence. The information-theoretic based metrics, Cosine Similarity and Document Disparity, are proposed as indicators of Relevance and Coherence for a piece of data. The proposed metrics are validated through a text-based experiment using Twitter data.

1 Introduction

Data nowadays are produced at an unprecedented rate; Cheap sensors, existing and synthesised datasets, the emerging internet of things, and especially social media, make the collection of vast and complicated datasets a relatively easy task. The era of big data is obviously here but unfortunately without an equal advance in the science of understanding it yet. With limited time and human power, the ability to effectively harness the power of big data is a problem encountered by many companies and organizations. Enormous amounts of data take up too much storage and computing resources. That, however, does not necessarily mean an increase of valuable information or better actionable items. Only small amounts of data may address questions due to noise, redundancy, and non-relevance.

To increase effectiveness in data storage and handling, companies and organizations are focusing on driving robust data analysis techniques, such as Robust Principal Component Analysis [1] and K-medoids [2][3], to extract insights from data and downsize storage by keeping only the relevant information. However, the general theme of "garbage in, garbage out" still applies, and with big data, the problem becomes even more pronounced. For example, no meaningful patterns would be recognized if the data itself doesn't contain much information, or even worse, phantom patterns will appear if the data is not relevant.

Metrics that can evaluate the sufficiency, effectiveness and value of the collected data will bring incalculable benefits on selecting valuable data sets to analyze. Not only will the unnecessary cost of collecting redundant data be reduced, but also the efficiency of obtaining insightful results will be improved. The Data Readiness Level (DRL) measurement was first proposed by the Laboratory for Analytic Sciences at NC Sate University as a means to quantify that relevance. As the name suggests, the goal of DRL is to measure the relative readiness/richness of data to answer specific questions by various techniques.

DRL should be a generic measure applied to a variety of data modalities, with inference/decision/answer to a question as a shared goal. Specifically, DRL should be a function of at least four properties of data: Noisiness, Believability, Relevance, and Coherence. In this paper, we focus on data sets comprised primarily of documents since this is one of the most common datasets available and the literature in the field is rich enough to exploit various techniques. We assume that a set of documents is given to help answer a specific question. That dataset may contain relevant and irrelevant information and it does not have to be structured. The goal is to help guide a data scientist with a limited access to the entirety of the data corpus or alternatively with a limited time to analyze it, to quantitatively assess the capacity of a subset of that corpus to answer the question. This will be achieved by developing a computational measure to reflect the readiness of available data to answer such a question.

This is, to the best of our knowledge, the first time this notion of "goodness of unstructured data" has been addressed in the context of data analytics for seeking answers to specific questions. We discuss here an unsupervised approach to computing DRL metrics on a wide variety of data. Moreover, we demonstrate its successful application to a collection of tweets. For computational efficiency and tractability, we use the assumption of "Bag-of-Words" (BOW) in the case study on tweets. In the BOW model, a text is represented

*Department of Electrical and Computer Engineering, North Carolina State University. Email: {hguan2, ahk}@ncsu.edu

†Department of Mathematics, Florida Polytechnic University. Email: agentimis@flpoly.org

‡Laboratory for Analytic Sciences. Email: keiserjm@gmail.com

as the collection of its words, disregarding grammar and word order but keeping multiplicity [4]. To that end, topic modeling [5], often used in text mining and natural language processing as a dimension reduction, is the adopted strategy in our work. The Latent Dirichlet Allocation approach [6] was selected, but any other approach (e.g. Latent Semantic Indexing [7], or Non-negative Matrix Factorization [8]) could have been used, just as well. Cosine Similarity and Document Disparity are computed to measure Relevance and Coherence properties for a specific text collection. The underlying assumption is that a relevant set of documents should have high similarity to the question and low internal disparity, indicating high Relevance and high Coherence respectively. Our formulation is fully data-driven, hence more "natural" and better suited to the underlying structure of the data, and the proposed metrics are relatively easy to compute.

The rest of this paper is organized as follows. In section 2, we discuss various required properties for a sound DRL formulation and several relevant factors in its definition. In section 3 we provide an overview of the mathematical, statistical and machine learning relevant tools we used. In sections 4 and 5, we describe the theoretical framework of our proposed method and the methodology for implementing the framework. For validation as well as illustrative purposes, we present a practical example on tweets in section 6. Some of the related work is described in the section 7 followed by a conclusion section 8.

2 Definition of DRL

DRL is a function on a data set, which has been subjected to a sequence of transformations towards answering a query. In that sense, it reflects a *degree of maturity* towards accomplishing a target task. As a metric, it may be evaluated at various stages of transformation of data, and will hence reflect the efficacy of each of the various stages/analytics, or in one shot for the entire flow, with a goal of refining and distilling the data more closely to successfully resolve the query.

As a valid and useful measure, DRL has to satisfy the following properties:

1. Easy computability. This ensures that DRL remains more advantageous than actually carrying out the information analysis.
2. Stability and continuity. This will imply that any two close formulations (data, question, analytics) will yield similar DRL values.
3. Scalability. This will safeguard its viability with increasing data size.
4. A maximally unsupervised property, but with the allowance of fine tuning the parameters.
5. A discriminative property, to make it useful for a data scientist in reaching a decision about the data at hand.

These properties are intrinsic to our proposed techniques and, up to minor adjustments, can be shown to be stable under perturbations/changes of the input data, in our case, documents. The mathematical development of our DRL will further unveil that the following have to be accounted for:

1. **Data.** Given the application scope of DRL, diverse data modalities may be addressed and any proposed metric or framework must address a variety of data. The data may be noisy and unstructured.
2. **Analytics.** Given the great diversity of analytics/transformations and their specificity to problems, DRL may be operating along a number of possible dimensions, including time, cost, and computational complexity.
3. **Objective Question (Question of Interest).** The nature of the proposed question may vary widely. Thus DRL must be flexible enough to accommodate binary, quantitative and/or probabilistic questions and answers. Additionally, questions may vary in substance and detail, and hence emanate from a population with a certain distribution, so an associated DRL will accordingly be in some sense a "conditional" quantity relative to the query.

To further clarify the problem statement, we focus on the application of DRL to unstructured text data. To that end, we assume that a corpus of documents is given to help answer a specific question. The goal of DRL is to quantitatively describe the amount of valuable information contained in the set of documents related to the question. We consider four dimensions of the set of documents in computing DRL:

1. **Relevance:** the extent to which data is applicable and helpful for the question at hand.
2. **Coherence:** the extent to which data is consistent in subjects or topics.
3. **Believability:** the extent to which data is regarded as true and credible.
4. **Noisiness:** the extent to which data is clean and correct.

The dimensions can be modified or discarded depending on the application. For example, for time-sensitive questions, the timeliness dimension should also be considered. In this paper, we account for only these four dimensions of data to make DRL as generic as possible.

Mathematically, let $\mathcal{D}$ be the space of documents and let D be a finite set of documents of interest with

cardinality N, i.e. $D = \{d_1, \cdots, d_N\} \subset \mathcal{D}$, where d_i corresponds to the i-th document. Similarly, let $\mathcal{Q}$ be the space of questions and let $\mathbf{q} \in \mathcal{Q}$ be a specific question. Our goal is to compute $DRL(D|\mathbf{q})$. The connection between the space of documents and that of questions is obvious if we treat questions as micro-documents. In this sense $\mathbf{q} \in \mathcal{Q} \subset \mathcal{D}$.

Ideally, DRL should be a function that maps any set of documents to some fixed numerical range, for example, $[0, 1]$, with 0 indicating no value at all and 1 indicating the most valuable data: $f : 2^{\mathcal{D}} \to [0, 1]$. An application of this would be to focus on the documents with the maximum possible DRL. Given the sets of documents $\{D_1, D_2, \cdots, D_M\}$, it is reasonable to only require the property: $DRL(D_i|\mathbf{q}) > DRL(D_j|\mathbf{q})$ i.e. D_i is more valuable than D_j on answering the question $\mathbf{q}$.

3 Background

To cast our proposed DRL methodology in an applied setting, and keep this report self-contained for a document-based experimental evaluation, we provide some brief background on text mining, making connections with the relevant research in the area.

3.1 Latent Dirichlet Allocation Since our DRL experimental validation involves document analysis, it is natural to approach documents as bags of words, as often done in classic data mining. The size of a given corpus of documents may, however, be prohibitively large making this approach computationally expensive and negatively impacting the efficiency of the DRL computation. To alleviate that problem we have adopted some tools from topic modeling and specifically the concept of Latent Dirichlet Allocation [6]. We assume thus that the question and the document that are addressing the same topics are likely to be relevant and use that as a bottom line for our relevance measure.

The essence of LDA is that documents are distributions over topics, the latter being themselves probability distributions over words. Suppose that $D = \{\mathbf{d}_1, \cdots, \mathbf{d}_N\}$ is a set of N documents. Let K be the number of topics and V be the number of words in the vocabulary. Let $\boldsymbol{\varphi}_k \in \mathbf{R}^V, k = 1, \cdots, K$ be the distribution of words corresponding to topic k. Given a document $\mathbf{d}_i \in D$, let N_i be the number of words and $\boldsymbol{\theta}_i \in \mathbb{R}^K$ be the distribution of topics. The generative probabilistic model defined in LDA is as follows:

1. Choose $\boldsymbol{\theta}_i \sim Dir(\boldsymbol{\alpha})$ $i \in \{1, \cdots, N\}$, where $Dir(\boldsymbol{\alpha})$ is the Dirichlet distribution with parameter $\boldsymbol{\alpha} = \{\alpha_1, \cdots, \alpha_K\}$.
2. Choose $\boldsymbol{\varphi}_k \sim Dir(\beta)$ $k \in \{1, \cdots, K\}$ where, $Dir(\beta)$ is the Dirichlet distribution with parameter β.
3. For each word $w_{i,j}$ in j^{th} position of the i^{th} document, where $i \in \{1, \cdots, N\}, j \in \{1, \cdots, N_i\}$:
   (a) Choose a topic $z_{i,j} \sim \text{Multinomial}(\boldsymbol{\theta}_i)$,
   (b) Choose a word $w_{i,j} \sim \text{Multinomial}(\boldsymbol{\varphi}_{z_{i,j}})$.

3.2 The Semantic Space of Queries To better contextualize DRL conditioned on a question $\mathbf{q}$, a class of questions Q conveying a similar "message" with possibly different words, is sought. This may be thought of as defining a semantic equivalence class to $\mathbf{q}$. Let again $\mathcal{D}$ be the set of all documents, each considered as a bag of words using a high dimensional lexicon. Each question, regarded as a mini document, would correspond then to a representation in the *Semantic Space* denoted by $\mathcal{S}$. This space may roughly be associated to the space of distributions over topics, which, recall, are the result of the LDA turning documents (a bag of words) into distributions over topics. Thus what we should be basing our DRL on is the set of questions which have the same image under some function: $g : \mathcal{D} \to \mathcal{S}$. The idea of this semantic space can be easily extended to other modalities, so long as a good proximity measure is identified. Unfortunately, such a measure is still fairly not well understood just like the notion of this space itself.

To proceed with the mathematical development, denote by D a family of documents which is a subset of $\mathcal{D}$. The LDA approach thus defines a function,

$$g_D : \mathcal{D} \to \mathcal{S}, \tag{3.1}$$

hence associating with each document its semantic meaning. A question $\mathbf{q}$ is typically also projected onto this Semantic Space through the same function g_D. By construction this function g_D depends on the document set itself, and hence the corresponding index D. It would be ideal if a general function could be found to attribute meaning to documents without the use of the specific setting; however, g_D is good enough for our purpose.

As shown in figure 1, the documents $\mathbf{d_i} \in D \subset \mathcal{D}, i = 1, 2, 3$ are projected onto the semantic space, and $\mathbf{q_1}, \mathbf{q_2}$ are two queries in $\mathcal{Q} \subset \mathcal{D}$, with the same representation in $\mathcal{S}$. The "distance" between the query and the documents should be computed in terms of distances in the semantic space. In so doing, the DRL will depend on the "meaning" of the question posed, and avoid being indicative of the specific words used to formulate it.

We note, that a philosophical analogy may be drawn between this formulation of the DRL with that of Information Retrieval(IR) as related to question expansion, which invokes techniques of finding synonyms, stemming, spelling correction, etc.. We maintain that DRL

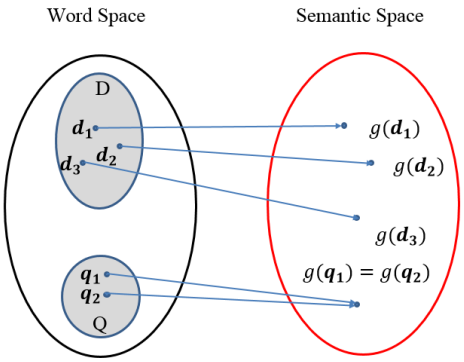

Figure 1: Space mapping

should, however, go further and deeper by breaking down the concept for better understanding of the question, and for discovering the intent behind the question.

3.3 Jensen-Renyi Divergence Information theory and statistics are rich with measures of closeness (conversely divergence) between probability density functions (PDF). On the other hand, a PDF, itself, is a reflection of intrinsic behavior of a given population. It then makes eminent sense to seek the same notion of the concentration of data in a given population, across populations. That is precisely what is reflected in information theoretic measures such as divergence (for example, Kullback-Liebler (KL) divergence). To that end, we build on some of our previous work in developing measures across an arbitrary number of PDFs , as a step beyond the well known and classical KL divergence between two PDFs. Viewing hence all documents as distributions over topics using the LDA technique, we proceed to first define the so-called Jensen-Renyi (JR) divergence [9].

DEFINITION 1. *Let $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_n$ be n probability distributions on a finite set $x_1, x_2, \ldots, x_k$, $k \in \mathbf{N}$. Each probability distribution is $\mathbf{p} = (p_1, p_2, \ldots, p_k)$, $\sum_{j=1}^{k} p_j = 1, p_j = P(x_j) \geq 0$. Let $\boldsymbol{\omega} = (\omega_1, \omega_2, \ldots, \omega_n)$ be a weight vector and $\sum_{i=1}^{n} \omega_i = 1, \omega_i \geq 0$. The JR divergence is defined as:*

(3.2)
$$JR_{\alpha}^{\boldsymbol{\omega}}(\mathbf{p}_1, \ldots, \mathbf{p}_n) = R_{\alpha}\left(\sum_{i=1}^{n} \omega_i \mathbf{p}_i\right) - \sum_{i=1}^{n} w_i R_{\alpha}(\mathbf{p}_i),$$

where $R_{\alpha}(\mathbf{p})$ is the Renyi entropy defined as:

$$R_{\alpha}(\mathbf{p}) = \frac{1}{1-\alpha} \log \sum_{j=1}^{k} p_j^{\alpha}, \alpha > 0, \alpha \neq 1.$$

The JR divergence is a convex function over $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_n$ for $\alpha \in (0, 1)$. It achieves a minimum value of zero when $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_n$ are equal for all $\alpha > 0$ and gets its maximum value when $\mathbf{p}_i = \delta_{ij}$, where $\delta_{ij} = 1$ if $i = j$ and 0 otherwise. Many more properties of JR divergence can be found in [9].

3.4 Sensitivity and Closeness In numerical linear algebra and linear system theory, a non-singular matrix A is ill-conditioned if a relative small change in A can cause a large relative change in $A_{-1} : (||x - x'||)/||x|| \leq k * ||B||/||A||$, where B is the perturbation input to the system A and k is the condition number $k = ||A|| * ||A^{-1}||$ [10]. This concept is extended to include non-linear systems, like the LDA, using ideas from functional analysis on metric spaces as follows:

DEFINITION 2. *Given a function $f : X \to Y$ where (X, d_x) and (Y, d_x) are metric spaces, we say that f is r-locally l-Lipschitz if and only if for all $x \in X$ and for all $y \in X$ such that $d_X(x, y) \leq r$ we have that:*

$$d_Y(f(x), f(y)) \leq l \cdot d_X(x, y).$$

The idea of a Lipschitz constant l is connected to the derivative of a function at a point (recall that a differential invokes a denominator as $d_X(x, y)$). If l is smaller than 1 we say that f is a contraction or in other words makes small errors even smaller (stabilizes). If the two spaces, however, have wildly different metrics, the Lipschitz constant is not very informative. To alleviate this problem we normalize the two metrics and define a new sensitivity number as follows:

DEFINITION 3. *(**Sensitivity Number**). Given a function $f : X \to Y$ where $(X, \|\cdot\|_X)$ and $(Y, \|\cdot\|_Y)$ are spaces with norms, we define the relative r-sensitivity at the point x_0 to be:*

$$s_1(x_0) = \inf\{\frac{\|f(x) - f(x_0)\|_Y}{\|f(x_0)\|_Y} \div \frac{\|x - x_0\|_X}{\|x_0\|_X} | \quad x \in X, \|x - x_0\|_X \leq r\}. \tag{3.3}$$

We will use these notions in Eq.3.3 to report a form of stability to small perturbations for our measures.

4 Theoretical Framework

In this section, we will define the two preliminary measures we propose for a functional DRL that correspond to the relevance of a set of documents to a certain question and its overall coherence.

DEFINITION 4. *(**Relevance** $\mathcal{R}$). Let M denote the number of sets of documents and N_i denote the number of document in the i-th collection. Each document $\mathbf{d}_j^{(i)}$ is represented as: $g(\boldsymbol{d}_j^{(i)}) = [\theta_{j1}^{(i)}, \theta_{j2}^{(i)}, \ldots, \theta_{jK}^{(i)}]$,*

$i = 1, 2, \ldots, M, j = 1, 2, \ldots, N_i$, *where* $\theta_{jk}^{(i)}$ *is the proportion of topic* k *in the* j*-th document on the* i*-th collection, using eq.(3.1). The question is also represented by* $g(\mathbf{q}) = [\theta_1, \theta_2, \ldots, \theta_K]$. *Define then a Relevance metric, to be the cosine similarity measure between each document in the collection* D_i *and the question* $\mathbf{q}$*:*

$$Sim(D_i, \mathbf{q}) = \frac{1}{N_i} \sum_{j=1}^{N_i} \frac{g(\mathbf{d}_j^{(i)}) \cdot g(\mathbf{q})}{\|g(\mathbf{d}_j^{(i)})\| \cdot \|g(\mathbf{q})\|}. \tag{4.4}$$

Formally, let $\mathcal{X}$ be a collection of sets of documents and $\mathcal{Q}$ be a set of questions, the function, $Sim : \mathcal{X} \times \mathcal{Q} \to \mathbf{R}^+$ is a DRL measure which denotes the relation between each set and a question. Note that this may just as well be applied to individual documents (when the document set size is 1) or to the whole corpus of documents. Eq.(4.4) obviously induces a relation between two document sets as follows:

Definition 5. *Let* D_1 *and* D_2 *be two sets of documents and* $\mathbf{q}$ *an associated question as described above. We say that* D_1 *and* D_2 *are informationally equivalent relative to the question* $\mathbf{q}$ *if and only if*

$$Sim(D_1, \boldsymbol{q}) = Sim(D_2, \boldsymbol{q}).$$

It is hard to find informationally equivalent document sets in practice, hence warranting a definition of δ-equivalent documents as follows:

Definition 6. *Let* D_1 *and* D_2 *be two sets of documents and* $\mathbf{q}$ *an associated question as described above. We say that* D_1 *and* D_2 *are* δ*-informationally equivalent relative to the question* $\mathbf{q}$*, if and only if*

$$|Sim(D_1, \boldsymbol{q}) - Sim(D_2, \boldsymbol{q})| \leq \delta.$$

If we consider the set $\mathcal{X}$ under the informational equivalence relation, or the δ informational equivalence, then the function Sim induces a total order relative to the question $\mathbf{q}$. Hence, all sets in $\mathcal{X}$ can be compared and the one which is most relevant to the question can be identified.

While cosine similarity measures which document set is more related to the question, it does not provide any information on how the document set is focused on the topic(s) of interest. It is possible that the document set with higher similarity with the question contains much irrelevant information. To explore that property we propose the idea of Document Disparity as an indicator of Coherence:

Definition 7. *(**Coherence** $\mathcal{C}$). Let* D_i *be a collection of documents written as distributions over topics. We then define the Document Disparity to be the JR divergence of this set of documents and the reciprocal of the document disparity will be the corresponding Coherence:*

$$\mathcal{C}(D_i) = \frac{1}{DD(D_i)} = \frac{1}{JR_{\alpha}^{\boldsymbol{\omega}}(g(\mathbf{d}_1^{(i)}), \ldots, g(\mathbf{d}_{N_i}^{(i)})))}. \tag{4.5}$$

This function, $\mathcal{C} : \mathcal{X} \to R^+$ is independent of the question $\mathbf{q}$, and measures how "different" the documents are within their set. A lower disparity measure within a set of documents is equivalent to saying that they are focused on the same category of topics, namely, more coherent.

In our experiments below we will prove that the metrics defined above are contractions and have a low sensitivity number as defined in section 3.4. In general, any relevant DRL should enjoy these properties.

5 Methodology

Given a set of document collections $\mathcal{X} = \{D_1, D_2, ..., D_M\}$, Figure 2 shows our framework to compute the aforementioned DRL metrics. Our method is comprised of two parts: the first involves a model training process, while the second evaluates the metrics for relevance and coherence using the previous definitions of relevance and coherence.

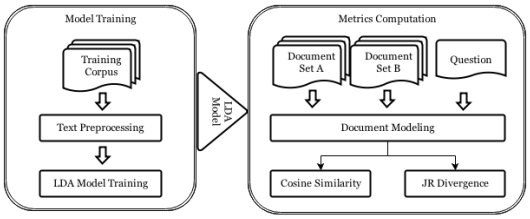

Figure 2: Methodology flow chart

In model training, a training corpus $D \subset \mathcal{X}$ is preprocessed to build an LDA model. Let V be the number of words in the vocabulary. According to the BOW assumption, each document $d \in D$ is a V-dimensional vector $d \in \mathbb{R}^V$, a point in the so-called "word space". The vector is sparse because only a few words in the vocabulary will appear in the document. The "LDA Model Training" module takes as inputs the vector representation of the training corpus D and the number of topics K, and generates an LDA model, $g_D : \mathcal{X} \to \mathcal{S}$, that could map a single document $\mathbf{d}$ from a "word space" of V dimension to a "semantic space" of K dimension. Note here that the training corpus D does not need to be the set of document collections $\mathcal{X}$. Other than the LDA model, other representation learning approaches such as doc2vec [11] and topic modeling methods [5] could be applied to approximate the mapping to the semantic space. Since we are more

interested in analyzing the results of DRL, we will leave the analysis of the best document vector representation for the future.

The second part of the process is the computation of the defined metrics as follows. Given a set of document collections $\mathcal{X}$ and a specific question $\mathbf{q}$, we project each document $\mathbf{d} \in D_i$, $D_i \in \mathcal{X}$, $i = 1, 2, \ldots, M$ and the question $\mathbf{q}$ to the semantic space $\mathcal{S}$ using the learned function g_D. Both cosine similarity $Sim(D_i, \mathbf{q})$ and coherence $C(D_i)$ are computed in the semantic space for all $D_i \in \mathcal{X}$ with Eq.4.4 and Eq.4.5 respectively.

6 Experiment and Results

6.1 Dataset Collection To validate our proposed DRL measures, we carried out the following experiments using Twitter data that we collected. Specifically, the dataset consists of a set of 463,790 tweets collected from Twitter's Streaming APIs[1] during the FIBA Basketball Word Cup 2014. The collected tweets cover a time span from August 30, 2014, to September 18, 2014.

To circumvent structural difficulties associated with Twitter Data when converting text into vectors, we performed the following pre-processing steps: stopwords such as "the","of" and "and" were deleted; HTTP links were deleted, leaving only digits and Latin characters; words of document frequency smaller than 5, and all 1 word-tweets were deleted; text was lowercased.

The cumulative number of 463,204 tweets yielded a vocabulary size of 12,207. The corpus of tweets was divided into daily sets resulting in a series $\mathcal{X} = \{D_1, D_2, ..., D_{20}\}$. The primary objective of the defined experiment was the discovery of the most relevant set to the following question: *Will the USA Basketball team win the world cup in Spain?*. We partitioned the tweets to match the FIBA basketball world cup daily schedule. The assumption underlying our experiment is there being an uptick in the messages of the day's game so that the tweet messages in the days approaching the final should strengthen the relevance of the set to the initial question of interest, namely the ultimate "win of the cup by the USA". All tweets about unimportant (or of least relevance) games on a given day are effectively noisier. Conversely, all Team USA games around key dates (preliminary rounds, knockout games etc.) lead to a large number of daily tweets, which in turn reflect the fans' opinion-prognosis about the question at hand.

6.2 Experiment Design Each tweet was regarded as a document; note that for simplicity, no tweet pooling schemes [12, 13] (aggregating tweets into one document) were used (albeit another viable alternative to consider). All the tweets were used as the training corpus to train the LDA model. After some experimentation, we converged on 50 as the number of topics of choice, as it gave a set of clear patterns; it is worth noting that this did not greatly impact the final results, but presented some computational advantage. This number is also a key parameter for the LDA algorithm, which for the most part represents the heaviest computational load in the experiment. For the actual implementation we used the Gensim Python Library [14] and run a Standard LDA on the corpus.

For the metrics computation, we first used LDA to project the collection of daily tweets as well as the question to the semantic space. We then computed both the *cosine similarity* between the tweets and the question (as distributions over topics) as well as the *Coherence* of each set. In order to account for the non-deterministic nature of LDA, we performed $n = 200$ random computations and have retained the average of the cosine similarity and JR divergence.

6.3 Experiment Results Our experiment results with the FIBA basketball world cup tweet data, are shown in Figures 4 and 5. We first note the emergence of a time-lag of at most one day in our results, which is likely due to the difference of game occurrence time and the tweet reaction time.

The question, "Will the USA Basketball team win the world cup in Spain?", was represented as a distribution over topics as shown in figure 3. The corresponding topics (top 5 words) were:

Topic 4 : 0.272*win + 0.119*rose + 0.102*derrick + 0.085*point + 0.083*performance;

Topic 19 : 0.343*world + 0.318*cup + 0.301*basketball + 0.008*group + 0.006*liked;

Topic 31 : 0.203*spain + 0.156*world + 0.066*worl + 0.060*celebration + 0.056*cup;

Topic 34 : 0.287*usa + 0.228*team + 0.211*world + 0.192*cup + 0.011*home.

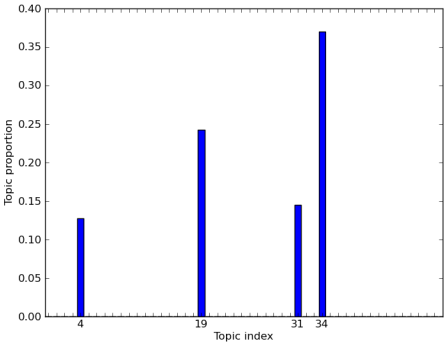

Figure 3: Topic proportion for the query

Figure 4 shows the average cosine similarity measure for the daily sets of tweets with the question $\mathbf{q}$, as

[1] https://stream.twitter.com/1.1/statuses/filter.json

well the variance over 200 iterations. We quickly note that the dates $9.2-9.3$, $9.11-9.12$ and $9.14-9.15$ show relatively high cosine similarity, with that of $9.14-9.15$ being the highest possible. According to the FIBA basketball calendar, USA qualified for the second round on 9.3, while mathematically on 9.2, thus explaining the high relevance of the tweets of that day to USA's team. We also note that on the 11^{th} of September, USA beat Lithuania during the semi-finals game and qualified for the final. The day of the final, 9.14, is clearly the day when the relation of the tweets is as close as possible to our question, thus revealing in the process the highest output that day.

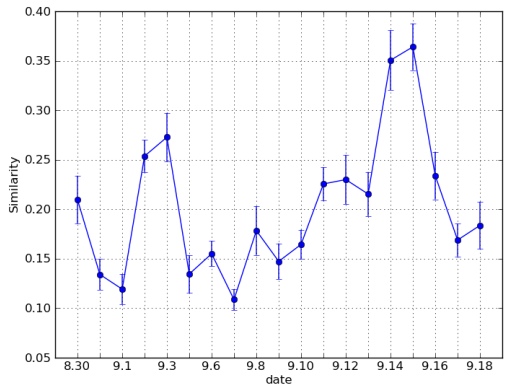

Figure 4: Daily Cosine Similarity

While one would expect a smaller spike in our graph on $9.9-9.10$, when USA beat Slovenia in the quarterfinals, another concurrent event (game of Lithuania versus Turkey) came in to spread the expected focus of the day on Team USA, caused by additional tweets about this other game. This amounts to declaring the dataset at that point "less ready" to answer the question of interest since more than one topic is discussed that day. This is more easily seen in the Document Disparity figure 5 where the DD for 9.9 is one of the highest recorded, so the coherence being the reciprocal of that is the lowest.

Looking now at figure 5, we can see that the lowest document disparity is found on $9.14-9.15$. Again this is easy to explain since the "Finals" is an event of central importance, and would hence be central to the tweets. We hasten to also point out that during the first few days $(8.30-9.2)$, a very high DD was due to the period of preliminary rounds when noise topics were emerging (at least one for each team participating). This trend carried on to $9.9-9.10$, at the conclusion of the quarter-finals, when the discussed topics are now more concentrated on the remaining teams and their upcoming final games. We believe the small dip in the disparity appearing on $9.2-9.3$ is related to the peak of the cosine similarity for those same days, and most likely on account of the fact that Team USA appeared set to qualify for the finals, and possibly a repeat victory. Various tweets of that day make predictions about the winners of the final and mention the fact that the USA is going to win again.

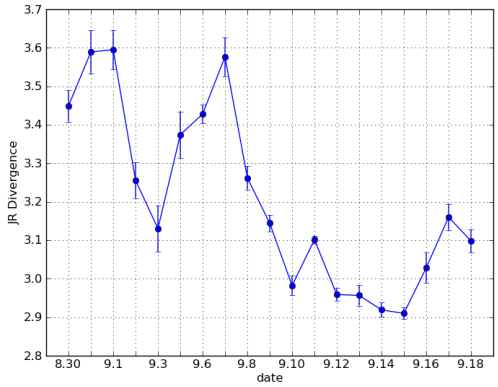

Figure 5: Daily Document Disparity

In light of the above results, we believe that both measures present a good foundation for DRL, by inferring relations of the data set to a question of interest. As a result, a data scientist may opt to bypass other sets of documents and focus on the ones with high cosine similarity and small document disparity. Conversely, one may argue, that for a certain class of questions, one should instead select a high cosine similarity and a high document disparity to reflect the highly exploratory nature of a posed question, rather that fact-driven. The proposed method, provides in both case, a step towards an automated, semi-supervised method to measure the readiness of a given data set. It remains that a true quantitative DRL is expected to be some *function* of the two metrics discussed above.

6.4 Perturbation Test Our goal is to show that any little perturbation of the query, $\mathbf{q}$, will only minimally impact the LDA model (same projection function) to hence provide a relatively consistent semantic representation, i.e. a small r-sensitivity number $s_1(\mathbf{q})$. Given the clear difficulty of computing the infimum over all possible perturbations of the query, we provide the sensitivity number for various different cases of perturbation and define a corresponding sensitivity quotient:

$$s_1(\mathbf{q}, \mathbf{q_p}) = \frac{\text{relative change in semantic representation}}{\text{relative change in query}} = \frac{||g_D(\mathbf{q}) - g_D(\mathbf{q_p})|| \cdot ||\mathbf{q}||}{||\mathbf{q} - \mathbf{q_p}|| \cdot ||g_D(\mathbf{q})||}.$$

In section 3 we defined the first DRL measure of a set of documents D with respect to a question $\mathbf{q}$ to be $Sim(D, \mathbf{q})$. The goal is to then see how sensitive the DRL measure is to a change in the query. Testing

this sensitivity number for a set of documents with high DRL(tweets from Sept.15) and another with low DRL (tweets from Sept.7) yields, much like $s_1(q, q_p)$, the normalized sensitivity quotient as,

$$s_2(X, \mathbf{q}, \mathbf{q_p}) = \frac{\text{relative change in DRL}}{\text{relative change in query}} = \frac{|Sim(X, \mathbf{q}) - Sim(X, \mathbf{q_p})| \cdot ||\mathbf{q}||}{||\mathbf{q} - \mathbf{q_p}|| \cdot |Sim(X, \mathbf{q}))|}.$$

We have constructed various perturbations of the query based on word repetition, synonyms replacement, and word deletion. Having fixed a dictionary, the bag of words representation of the query is similarly independent of the set of documents $\mathcal{X}$, all the while using the LDA model. Note though that the semantic representation through the LDA, heavily depends on $\mathcal{X}$.

Table 1: Perturbed query examples

Repetition
q_{a1}: USA USA basketball team win world cup Spain
q_{a2}: USA basketball basketball team win world cup Spain
q_{a3}: USA basketball team team win world cup Spain
q_{a4}: USA basketball team win win world cup Spain
q_{a5}: USA basketball team win world world cup Spain
q_{a6}: USA basketball team win world cup cup Spain
q_{a7}: USA basketball team win world cup Spain Spain
Replacement
q_{b1}: USA basketball team win FIBA Spain
q_{b2}: USA basketball team win world cup FIBA
q_{b3}: USA basketball team win world cup Spain2014
Deletion
q_{c1}: USA team win world cup Spain
q_{c2}: USA basketball win world cup Spain
q_{c3}: USA basketball team win world cup
q_{c4}: USA basketball team win Spain

A list of perturbed queries is shown in table 1. The topic distributions of some of these perturbed queries are shown in Figure 6. The sensitivity numbers are shown in Figure 7: The entries $1-7$ correspond to the queries $q_{a1} - q_{a7}$. The entries $8-10$ correspond to the queries $q_{b1} - q_{b3}$ and the entries $11-14$ to the queries $q_{c1} - q_{c4}$.

In almost all cases, the sensitivity quotient number s_1, corresponding to "LDA" in the figure, is smaller than 1. Only for the cases q_{c1} and q_{c2} is it bigger than one. This is because topics 19 (about the world cup basketball) and 34 (about the USA basketball team) were ill-proportioned in the perturbed query due to loss of information. One can suggest that this change (deletion of one word) is small in terms of distance in $\mathcal{W}$ but relatively large in the semantic space $\mathcal{S}$. The sensitivity quotient number s_2 is also smaller than 1.

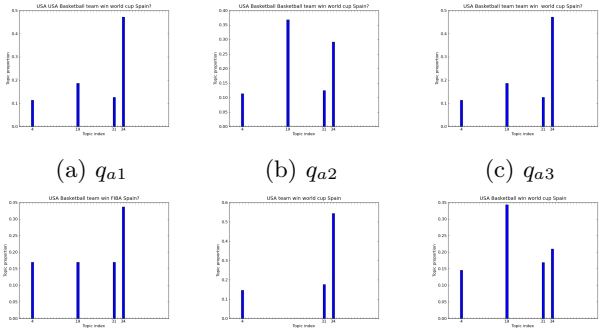

(a) q_{a1} (b) q_{a2} (c) q_{a3}

(d) q_{b1} (e) q_{c1} (f) q_{c2}

Figure 6: Topic Distributions of Perturbed queries

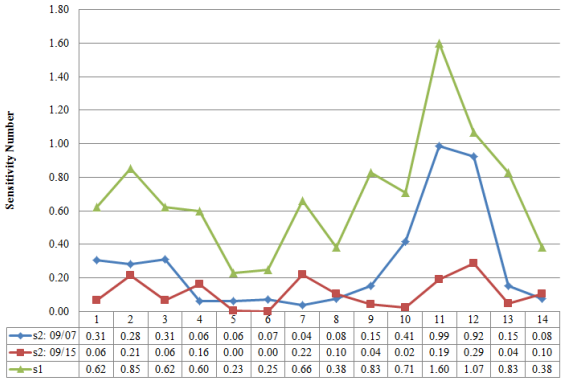

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
s2: 09/07	0.31	0.28	0.31	0.06	0.06	0.07	0.04	0.08	0.15	0.41	0.99	0.92	0.15	0.08
s2: 09/15	0.06	0.21	0.06	0.16	0.00	0.00	0.22	0.10	0.04	0.02	0.19	0.29	0.04	0.10
s1	0.62	0.85	0.62	0.60	0.23	0.25	0.66	0.38	0.83	0.71	1.60	1.07	0.83	0.38

Figure 7: Sensitivity numbers

7 Related Work

The idea of Data Readiness Levels is closely related to the field of information retrieval. The main goal of traditional information retrieval is to find relevant documents to a natural language query [15]. The desired output there is for the most relevant documents to be top-ranked in the returned list of documents. In our research, though, we would like to account for other data modalities such as videos and images. The general definition of DRL would evaluate the value of a piece of data to an objective while information retrieval measures only the relevance of a piece of information to a query. It is true that retrieved information based on the query should have a higher DRL than a randomly selected piece of data with respect to the same query. In this case, information retrieval techniques, especially question answering, could be applied to increase the DRL of data in the relevance dimension and probably coherence dimension. In cases where text is the main data medium, DRL is a different problem than information retrieval in that DRL focuses on the "goodness" of a collection of documents to an objective instead of a single document.

Technology Readiness Level (TRL) can be thought of as a template for DRL. TRL is a well-established

methodology of estimating the technology maturity during the acquisition process to assist decision-making concerning technology funding and transition [16]. The U.S. Department of Defense (DoD) has defined TRL based on a scale from 1 to 9 with 9 being the most mature technology. Unlike TRL, there are no standard evaluation rules for data maturity but we propose that it is not possible to define a general DRL without a specific objective. For example, a collection of documents about basketball games has less value to answer a question related to volleyball than to basketball no matter how refined these documents are.

8 Conclusions

In summary, cosine similarity measures the average relation of a document set and a given question, while document disparity reflects the variance or diversity of the information contained in a document set. Given the two measures, we would favor a data set with larger Sim measure and lower DD measure. Such a data set would have a high data readiness level.

By a combined theoretical-experimental approach, we have laid out some groundwork for a viable definition of a computable DRL measure, usable for any analytic process, especially for large data sets. While heavily intuited and illustrated using document-based data, DRL should be a generic and flexible measure compatible with any data modality. A great deal of work clearly remains to be undertaken, particularly in accurately establishing a quantitative scale for it, as well as applying to other modalities such as images, audio signals etc.

We believe that by exploring the idea of the semantic space, a concept which is not very well understood, the linking between various modalities and a specific question will be achieved, opening the way for comparison and computations on readiness. Here we are not using the standard ideas behind the semantic space as the re-formulation of documents through a topic discovery model. What we really mean is the creation of an abstract space, containing cognitive information and associations between ideas, similar to human intellect and understanding. In this semantic space, the concepts contained in documents, images, and other signals, will have a common reference system enabling us to link and compare them. Distances in this semantic space will formalize the idea of "notionally close" that humans naturally possess.

Acknowledgment

We would like to thank all the members of the Data Readiness Level Team especially Dr. Harish Chintakunta, for meaningful discussions on the subject, and the Laboratory for Analytics Sciences for its generous financial support.

References

[1] Emmanuel J Candès, Xiaodong Li, Yi Ma, and John Wright. Robust principal component analysis? *Journal of the ACM (JACM)*, 58(3):11, 2011.

[2] Leonard Kaufman and Peter Rousseeuw. *Clustering by means of medoids*. North-Holland, 1987.

[3] Brendan J Frey and Delbert Dueck. Clustering by passing messages between data points. *science*, 315(5814):972–976, 2007.

[4] Zellig S Harris. Distributional structure. In *Papers in structural and transformational linguistics*, pages 775–794. Springer, 1970.

[5] David M Blei. Probabilistic topic models. *Communications of the ACM*, 55(4):77–84, 2012.

[6] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3:993–1022, March 2003.

[7] Thomas Hofmann. Probabilistic latent semantic indexing. In *Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval*, pages 50–57. ACM, 1999.

[8] Daniel D Lee and H Sebastian Seung. Algorithms for non-negative matrix factorization. In *Advances in neural information processing systems*, pages 556–562, 2001.

[9] Yun He, A. Ben Hamza, and Hamid Krim. A generalized divergence measure for robust image registration. *IEEE Transactions on Signal Processing*, 51:1211–1220, 2003.

[10] David A Belsley, Edwin Kuh, and Roy E Welsch. *Regression diagnostics: Identifying influential data and sources of collinearity*, volume 571. John Wiley & Sons, 2005.

[11] Quoc V Le and Tomas Mikolov. Distributed representations of sentences and documents. In *ICML*, volume 14, pages 1188–1196, 2014.

[12] Rishabh Mehrotra, Scott Sanner, Wray Buntine, and Lexing Xie. Improving lda topic models for microblogs via tweet pooling and automatic labeling. In *The 36th Annual ACM SIGIR Conference*, page 4, Dublin, Ireland, July 2013.

[13] Liangjie Hong and Brian D. Davison. Empirical study of topic modeling in twitter. In *Proceedings of the First Workshop on Social Media Analytics*, SOMA '10, pages 80–88, New York, NY, USA, 2010. ACM.

[14] Radim Řehůřek and Petr Sojka. Software Framework for Topic Modelling with Large Corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, pages 45–50. ELRA, May 2010.

[15] ChengXiang Zhai. Statistical language models for information retrieval. *Synthesis Lectures on Human Language Technologies*, 1(1):1–141, 2008.

[16] US DoD. Technology readiness assessment (tra) guidance. *Revision posted*, 13, 2011.